

CodeBeam SF - 2018

BUILDING AND INTEGRATING A DATA PLATFORM



- benoît chesneau
- craftsman working on P2P and custom data endpoints solution
- **enki multimedia**: the corporate interface
- member of the **Erlang Industrial User Group**

about me

mieremobile services

1. Does my service do only one thing?
2. Is my service autonomous?
3. Does this service own its own data?

a good micro-service?

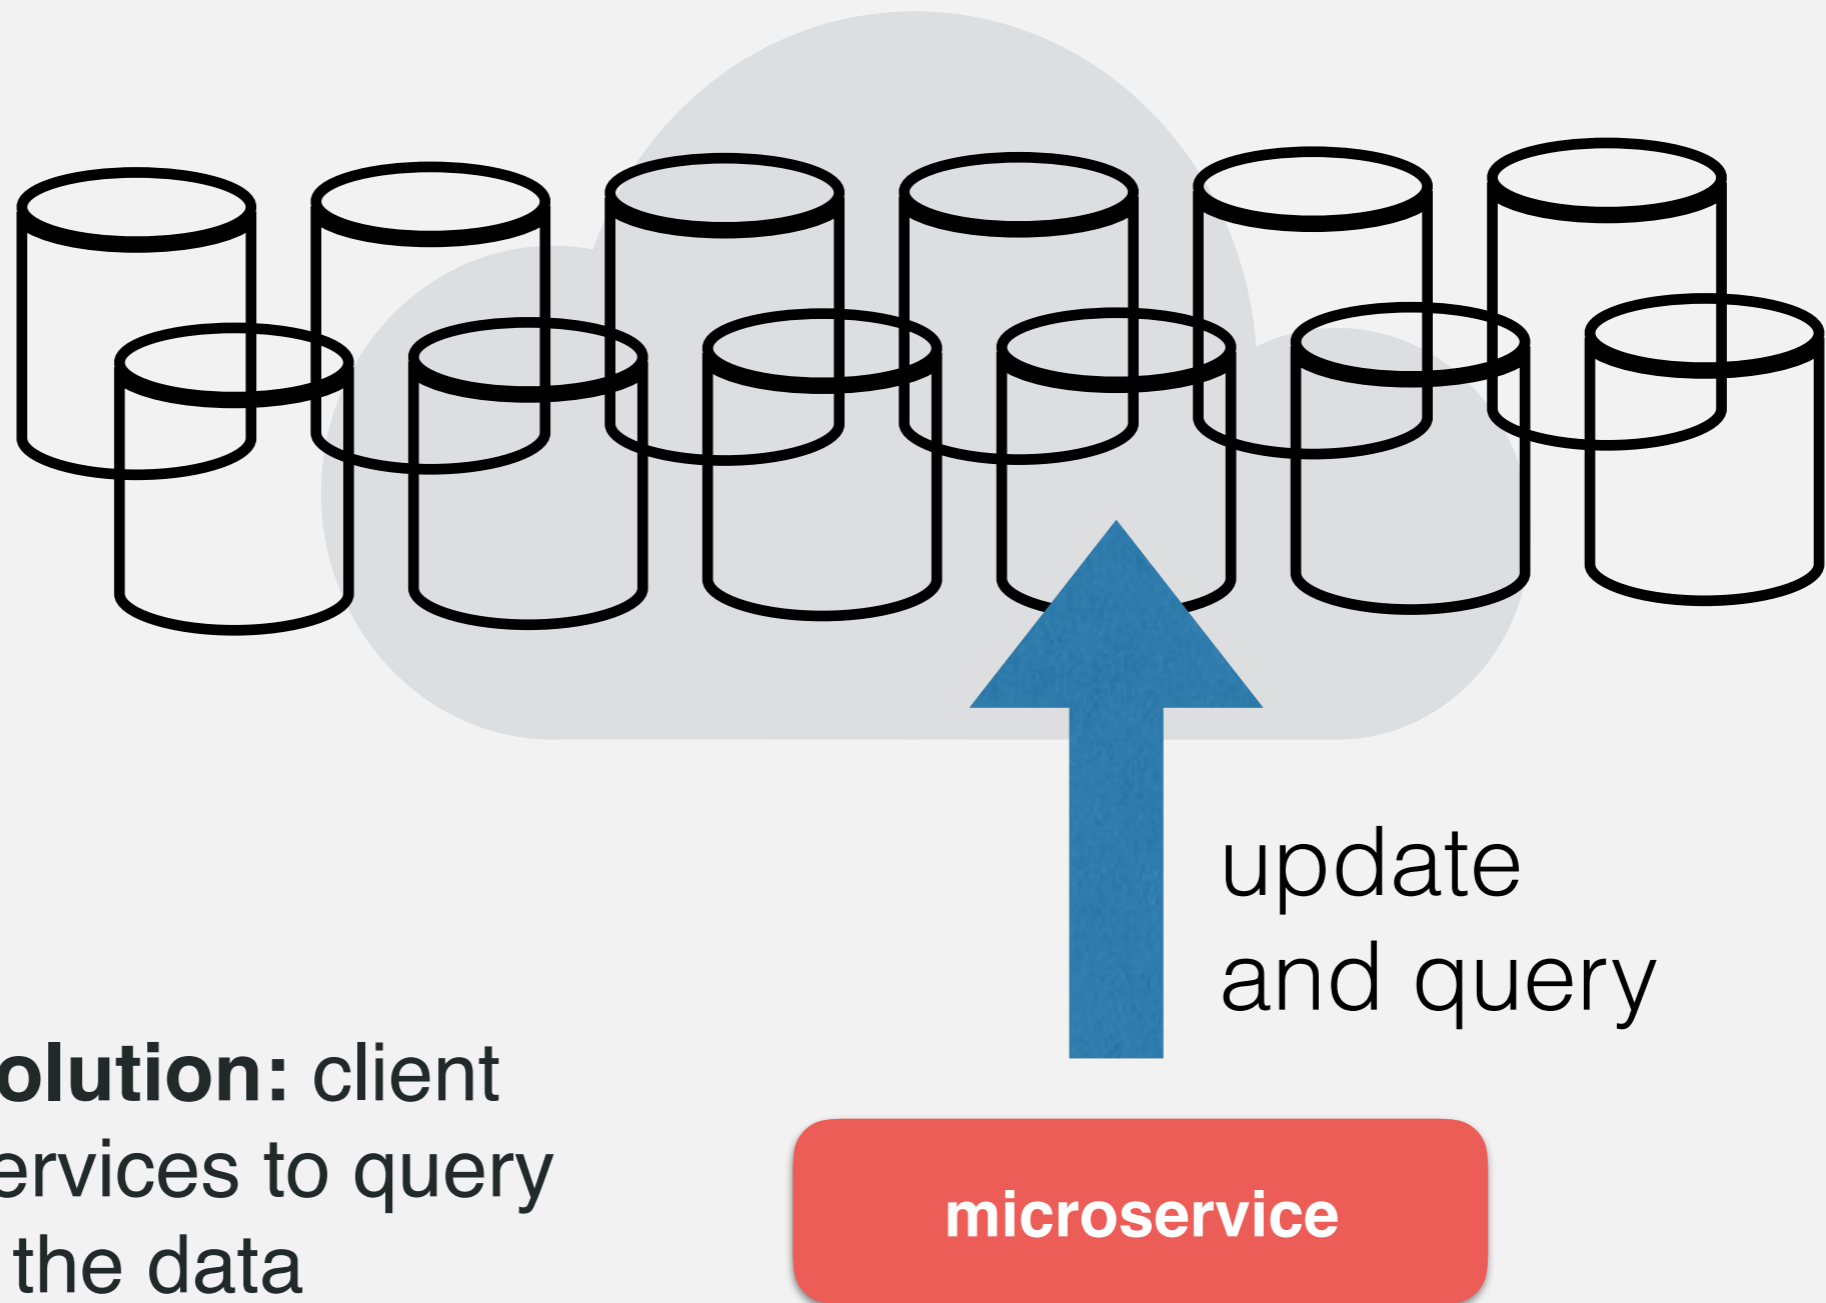
- isolated
- own its own data
- resilient
- communicate with other by asynchronous messages

micro-service

- sharing data in the mobile age between and across micro-services make applications more scalable and resilient
- Ex: messaging systems,

sharing data


cloud
storage



- **standard solution:** client call a webservice to query and update the data
- **problem:** if connection is slow or absent the microservice stops

sharing data

cloud
storage




synchronize



- **local storage** replicated
always available
- eventually consistent


local
storage



update
and query



microservice



sharing data

barrel

**Bring and keep a
view of your data
near your application**

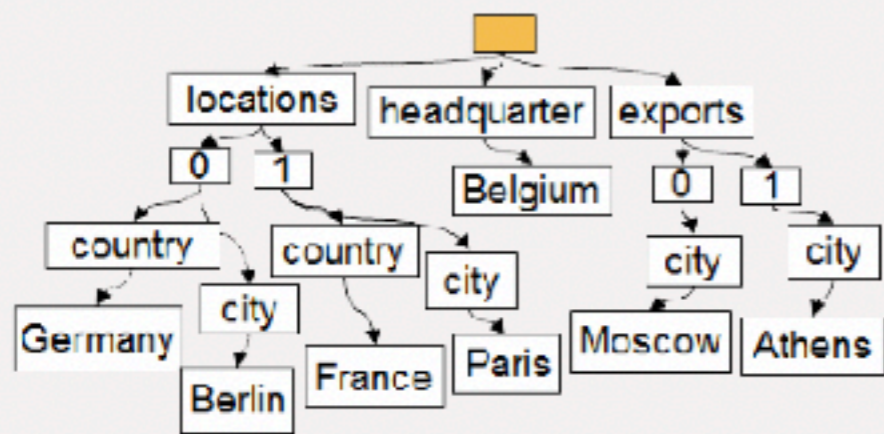
- a database focusing on simplicity
- document oriented
- Automatic indexing

Focusing on simplicity

```
{ "id" : "someid",  
  "Key" : "value" }
```

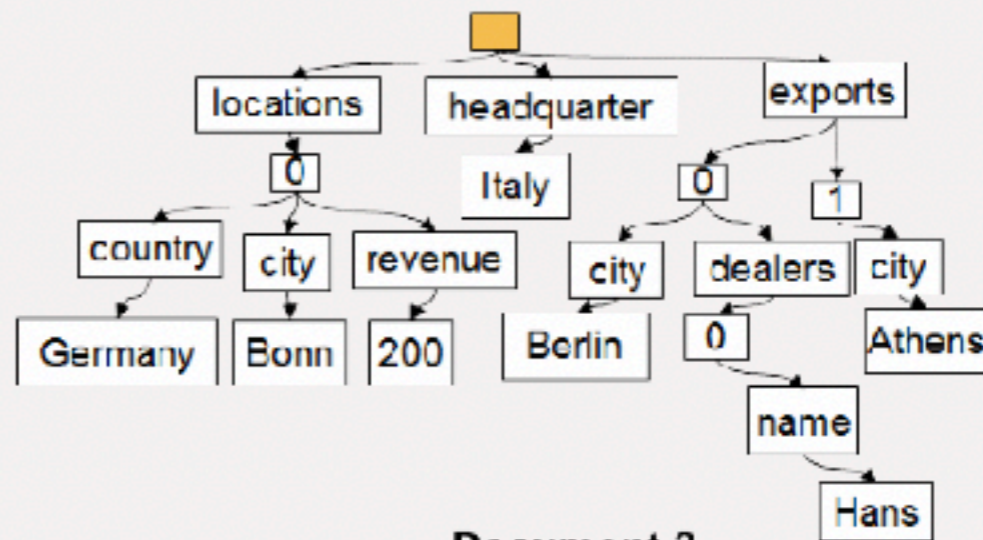
Docs are maps

```
{ "locations": [
  { "country": "Germany", "city": "Berlin" },
  { "country": "France", "city": "Paris" }
],
"headquarter": "Belgium",
"exports": [{ "city": "Moscow" },
            { "city": "Athens" }
]
};
```



Document 1

```
{ "locations": [
  { "country": "Germany",
    "city": "Bonn", "revenue": 200
  }
],
"headquarter": "Italy",
"exports": [
  { "city": "Berlin", "dealers": [{"name": "Hans"}] },
  { "city": "Athens" }
]
};
```



Document 2

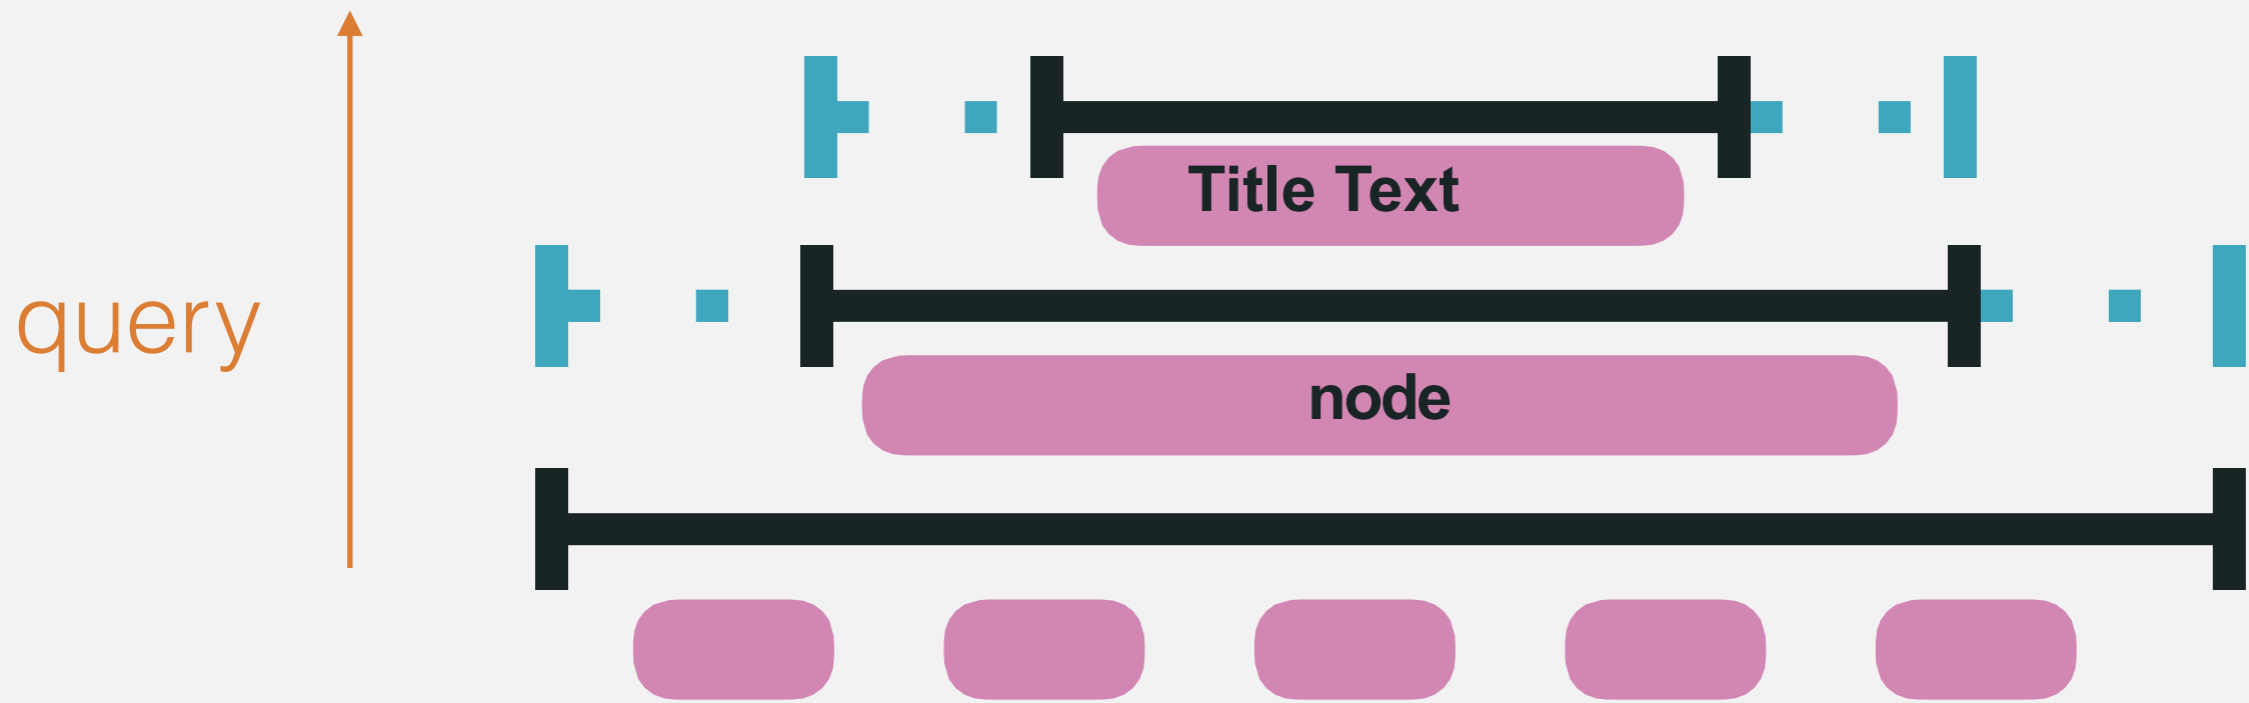
Access by path: /locations/country/Germany

automatic indexing



- local first: bring and keep a view of your data near your application
- data is synchronised with other storages
- Replication to and from any sources

Local first




partial view

- library embedded in your Erlang application(*)
- available as a micro-service via HTTP(1,2) or via the Erlang distribution
- Peer to peer: a barrel is the unit
- Semantic to allow distributed transactions

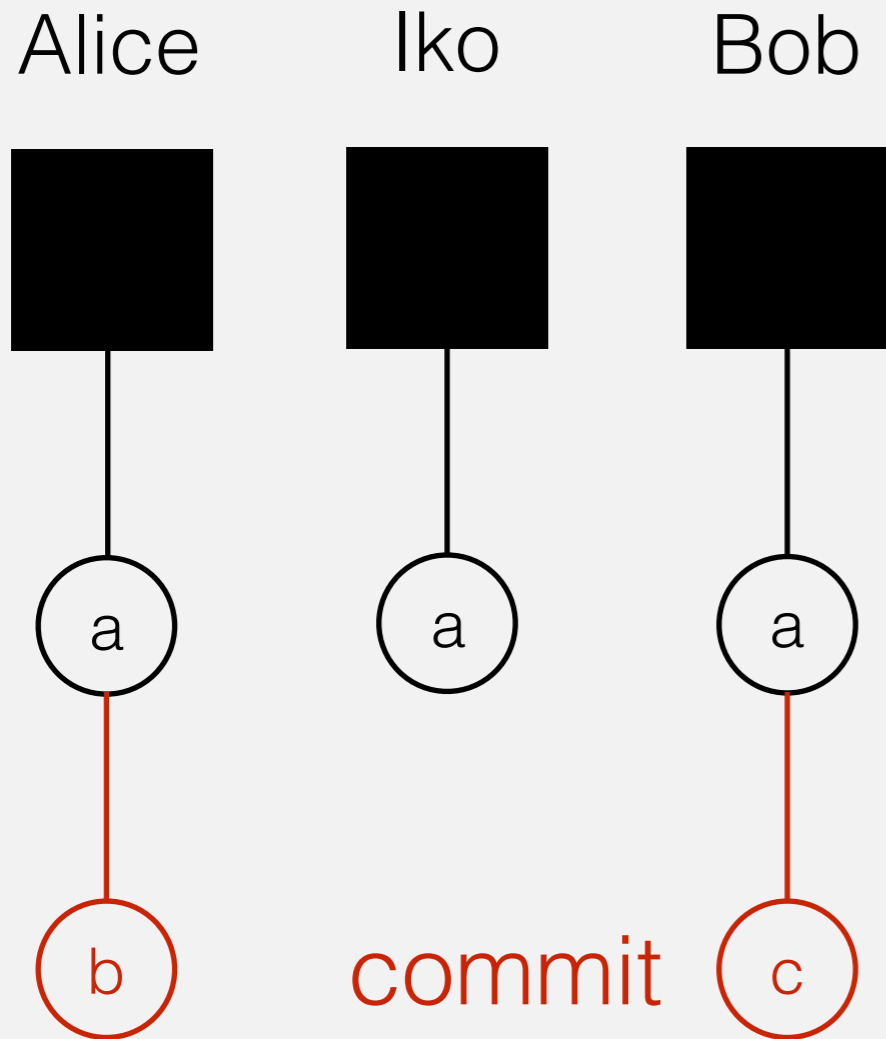
() including elixir or lfe, or*

- every peers fork the master, updates are offline
- peers pull and merge from the main server
- works well for back pressure (writes can be delayed)
- CRDT semantic for conflict-free data structures

- no vector clock
- revision tree

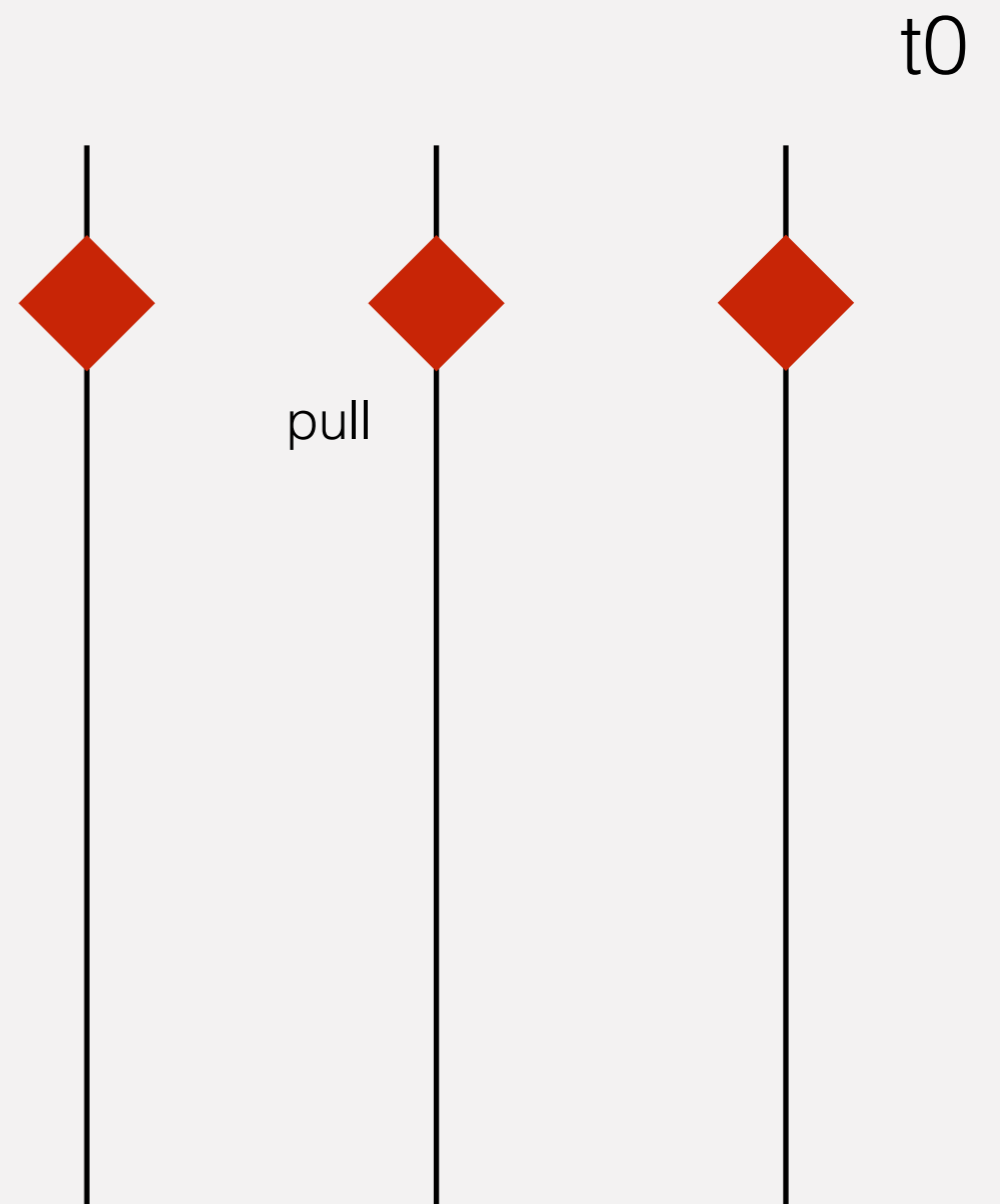


state




- ◆ manual
- ◆ automatic
- ◇ rejected

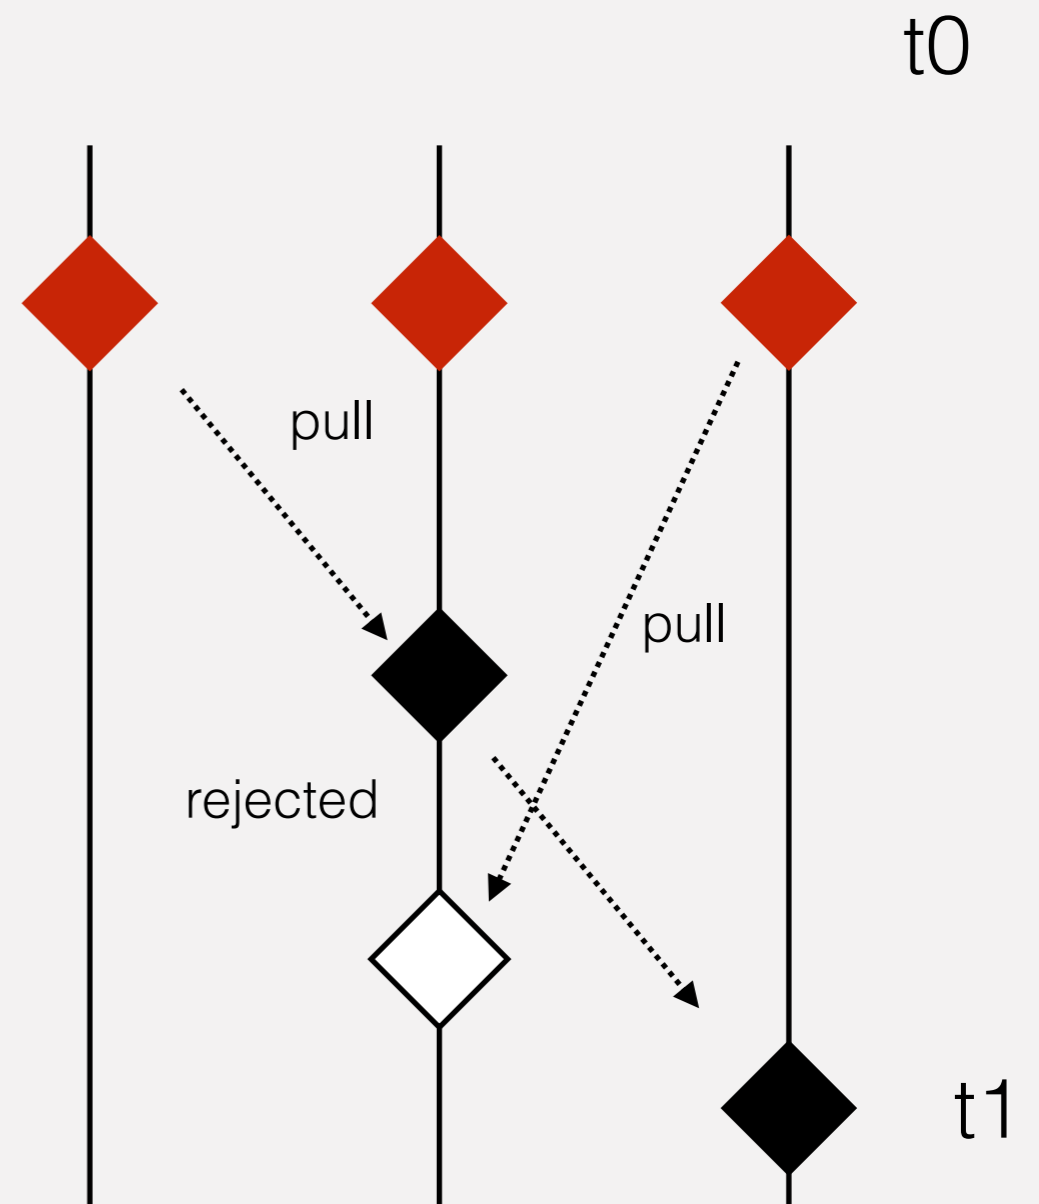
operations




state



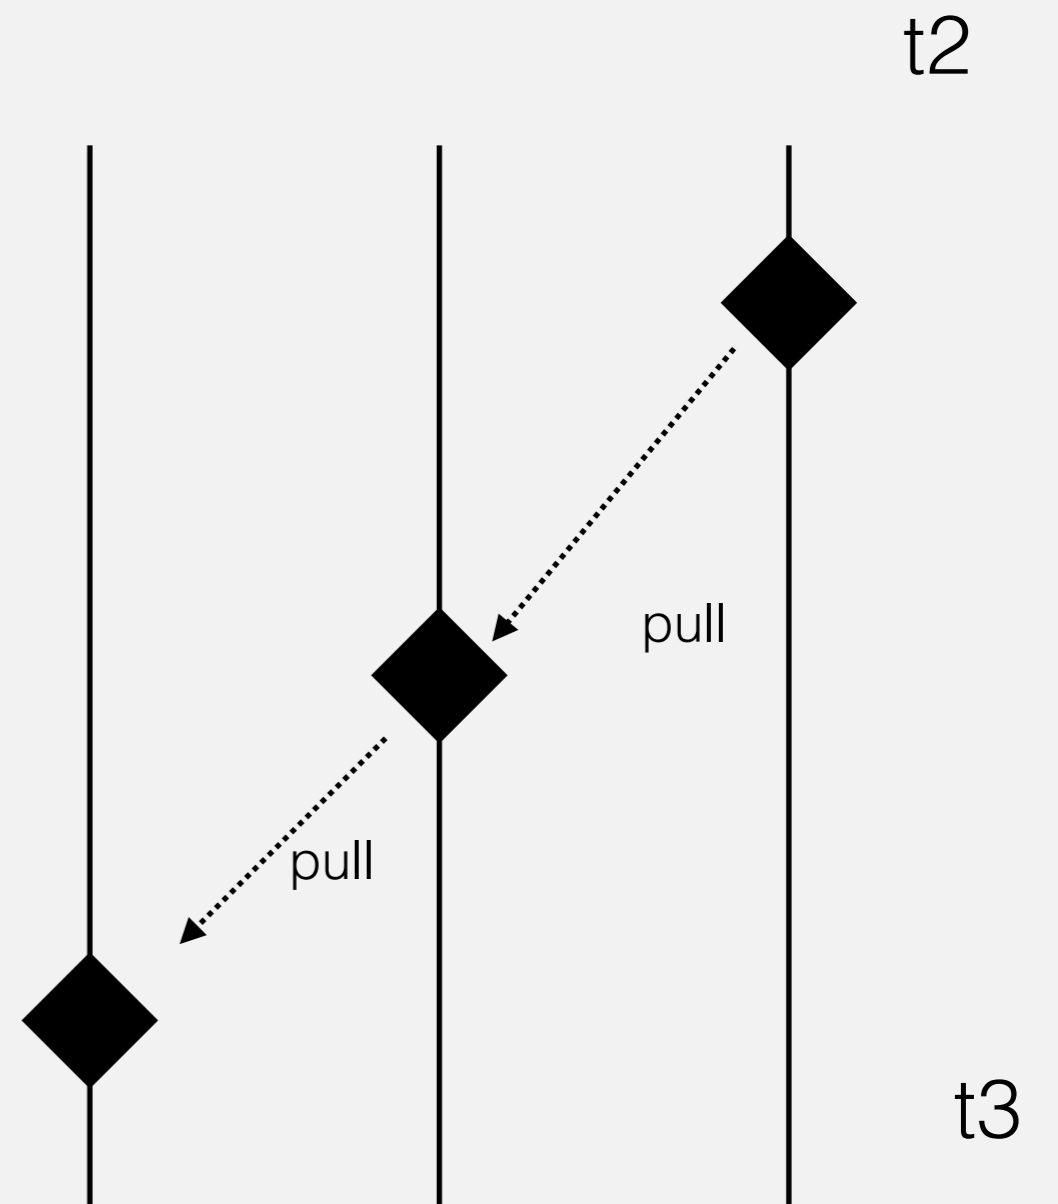
operations



state



operations



Erlang ???

- Erlang is slow
- Erlang is only for communications protocols
- I should do it rust...
- No access to low level memory and file systems
APIS

Why not Erlang

- Barrel is more a data orchestration service than a database
 - Basic indexing
 - Focus on replicating the data
- Nifs to help

Why Erlang

- Doc: Revision + Metadata data:
- Read-Modify-Write: concurrency issue
- Incremental changes log: append only
- Indexes: when a new winning version is found the doc is indexed.
- Blobs (attachments)

What we write

- Provides connectors for other storages
- Rocksdb for local persistent storage
<https://gitlab.com/barrel-db/erlang-rocksdb.git>
- Dirty-nifs
- ETS?

Use the right tool for ...

- Goal: anticipate the resource usage at the node level
- Return early to the client
- Control applied to all resources in the nodes
- Back-pressure

let it bend: be resilient

- worker_pool
https://github.com/inaka/worker_pool
- Hard to debug your program
- Little control on the pending requests
- Ecpcoxy but handle back-pressure the reverse way

Simple pooling

- Clients and Jobs should be handled independently
- Active and passive regularion
- Request unit: to set the number of requests we want to serve / seconds
- Flow-Based programming?
- sbroker, partially fit the bill:
<https://github.com/fishcakez/sbroker>

Dynamic regulation

- Started with a simple “Single Writer Multiple Readers” pattern
- bottleneck: A process to handle the final write to the database
- We do and // most of the work out of the write process
- Indexes are processed asynchronously (but a session can read its own writes if needed)

Concurrency challenge

- Read access is shared via ETS
- On request a monitor to the db is created

**ets: to share the state
between readers**

- When using the erlang distribution, events are dispatched by nodes, processes always subscribe locally

Events

- Erlang distribution is not used to share the data
- Erlang distribution can be switched
- HTTP transports

Transport the data

Roadmap

- 1.0: 24 march 2018
- 1.1: 24 april 2018
- ...

Milestones

- 1.0: Websockets support (with new hackney)
- 1.1: Experimental: GRPC

Coming features

?

- barrel is released in **march 2018**
- <https://barrel-db.org>
- contact me @benoitc